Re-Configurable Multimode, Multiband Information Transfer Systems

(Low Power, Small Size, Software Defined Radio, Secure Re-Configurable Mobile Communications & Computing)

M. J. Flynn, M. Morf
Computer Systems Laboratory
Stanford University
Stanford, CA 94305

Final Report, May 1999

20000522 037

Principal Investigator: Michael J. Flynn
Address: E.E. Dept., Gates Computer Science Building, Room 334, Stanford, CA 94305,
e-mail: flynn@ee.Stanford.edu , Phone: (650) 723-1450, Fax: (650) 725-6949

Investigator: Martin Morf
Address: E.E. Dept., Gates Computer Science Building, Room 335, Stanford, CA 94305,
e-mail: morf@arithmetic.Stanford.edu, Phone: (650) 723-0140, Fax: (650) 725-6949

Subcontractor: Mark Cummings, enVia,
Address: 348 Camino al Lago Atherton, CA. 94027, e-mail: cummings@envia.com,
Phone/Fax: (650) 854-4406

Final Report Table of Content

Project Summary

Our research was aimed at two Darpa programs, algorithms and reconfigurable architectures as part of the Adaptive Computing Systems (ACS) program, and applications to wireless networks, including coding and security as part of the Global Mobile Information Systems (GloMo) program.

Problem Statement:

We focused on two research problem domains: algorithms and reconfigurable architectures as part of Darpa's Adaptive Computing Systems (ACS) program, and applications to wireless networks, including coding and security as part of Darpa's Global Mobile Information Systems (GloMo) program. Wireless networks are ideal test domains for reconfigurable architectures.

Tomorrow's wireless networks can be characterized by an increasing number of transmitters and receivers and larger and larger data rates, and a growing profusion of modes of operation waveforms. The mobile communications user is confronted with a myriad of incompatible communications services. Each has its own technical, geographic, and feature set. A mobile defense professional moving through a mission moves in and out of different coverage areas; hence their communication needs change constantly.

Users need a *single communication device* that communicates with anyone, anywhere, with any available infrastructure. Therefore, one of the goals of this project was to research the necessary cost effective technologies that enable such devices.

Project Goals:

The *proposed goal of the research* was to understand and verify the full extent of the opportunities to satisfy secure defense mobile communications requirements for multi mode multi band hand held units with reconfigurable computing architectures. Milestones were based on the analysis, simulation, and implementation on a reconfigurable hardware test bed of new architectures and algorithms surrounding the addition of a widely deployed mobile infrastructure.

Research Summary:

Stanford's reconfigurable architecture research in ACS was based on its PamBlox approach to FPGA programming, an object oriented design methodology which has shown some significant advantages. Various arithmetic and security algorithms were implemented in efficient low-power FPGA designs, including DES and IDEA encryption standards.

Stanford also investigated new space-time code diversity based algorithms for wireless communication in the context of software defined radios, with the goal of providing or extending simultaneous *security* (privacy, integrity, reliability, availability, and protection from reverse engineering) by exploiting parallelism and scalable hierarchical multiplexing schemes to build software defined radio architectures. Stanford developed new Space-Time-Code Diversity (STCD) based *security architectures* that support all these requirements and map well onto software defined radio technologies.

Stanford demonstrated, as proposed, a multi-path approach to secure wireless data transmission in 1998 using three laptop computers, based on wired, modem, IR or Ethernet links, and commercial wireless Ricochet or Freewave links.

Morphics' goal was to show that multiple digital standards can be implemented in a single baseband architecture without multiple, parallel processing paths. The functional target was a multi-standard hardware logic section -- the new, previously undeveloped part of a multi-

standard communications baseband, to enable a *multi-standard digital modem*. The services considered during the project focused on TDMA based services.

Morphics presented two demos: a GSM FPGA baseband section, and a software structural design for a second service. The GSM standards were implemented and the IS-54/136 was in progress. A test bed was assembled based on Lucent's GSM test platform (Sceptre).

Morphics demo showed their completed design and demonstrated it on their Hardware Platform, GSM together with the analysis of IS-136 (North American TDMA), and IS-95 (CDMA), indicate that a multi-mode communication device is possible using reconfigurable architecture in combination with instruction-set processor(s) for baseband digital processing. The Morphics development at time of conclusion of the contract successfully shows that a complete range of TDMA standards can be implemented. The next stage would be to add the CDMA capabilities.

enVia's work focused on the RF frontend of a multiple service prototype system supporting GSM and TDMA & AMPS operating at cellular, 800-900 MHz, and PCS, 1800-1900 MHz frequencies. The current implementation involved LSI scale discrete analog components. VLSI single chip solutions are beginning to exhibit acceptable performance. Remaining discrete components are SAW filters, Dielectric Resonator Filters, Crystals. Digital/RF interfaces between the final Demod stage in the RF frontend and the high speed reconfigurable DSP backend include: Analog, Baseband, and I&Q. Control signals into the RF frontend are analog. This range of services requires three different interface bandwidths.

enVia **demonstrated** an RF frontend supporting 4 modes, 2 bands, and 3 bandwidths, AMPS, IS-136, IS54A, IS54B, IS54C, IS-95, PCS, PCS1900, IS136+, and IS-95+ on 5 RF test boards: Receiver, Transmitter, #1 LO, IF, Combiner Switchboard and Auxiliary LO.

enVia's RF frontend design involved a mix of discreet components to meet the project goals for performance, power efficiency, size, and cost. For IS-136, GSM, and IS-95, three IF SAW filters were used with frequencies in the 110 & 211 MHz bands, for package size and cost.

1. Problem Statement

We focused on two research problem domains: *algorithms and reconfigurable architectures* as part of Darpa's Adaptive Computing Systems (ACS) program, and *applications to wireless networks*, including coding and security as part of Darpa's Global Mobile Information Systems (GloMo) program. Wireless networks are ideal test domains for reconfigurable architectures.

Wireless Network Applications

Tomorrow's wireless networks can be characterized by an increasing number of transmitters and receivers and larger and larger data rates, and a growing profusion of modes of operation wave forms (Mod/Demod, RF IF & Baseband bandwidths, coding techniques, security, etc.)

The mobile communications user is confronted with a myriad of incompatible communications services. Each has its own technical (frequency, modulation, protocol, etc.), geographic (coverage, etc.), and feature set (voice, paging, data, etc.). As the mobile defense professional moves through a mission, they move in and out of different coverage areas and their communication needs change constantly.

In the US today there are at least four major cellular standards (AMPS, IS-54, IS-95, CDPD), and six PCS standards. One way paging is being expanded with two way paging. There are a growing number of wireless packet data services. In Europe, there are four similar standards (GSM, CT-2, DECT, and DCS1800) and more to come. In Japan, there are versions of these plus Handiphone. The number of services is clearly increasing. There are at least four Low Earth Orbiting Satellite Systems (LEOS) under development. Wireless LAN products are becoming more available. Wireless PBX's have arrived. This same process is going on in Europe, Asia and other parts of the world. For geopolitical reasons there are likely to be at least three non-compatible super sets of services. During this period of innovation, the number of incompatible services is likely to increase.

It is confusing, expensive, inconvenient and dangerous to deal with this jungle of services, with a single device for each service. Current technology solutions involve expensive, not mission adequate systems which can produce logistical nightmares such as requiring personnel to carry more weight in batteries than in food -- an incentive for foot soldiers to leave equipment behind. We note that during Desert Storm at least five generations of wireless communication systems were simultaneously active, and expected to cooperate!

What users want is a single device that will allow them to communicate with whomever they need to, wherever they are, with whatever infrastructure happens to be available. One of the major goals of this project was to research cost effective technologies that enable such devices.

Current DSP ASIC solutions suffer from problems with size and power consumption. Most systems employ an ASIC chip set for each communications service. These chip sets typically include a number of ASICs per service. The goal is clearly to design reconfigurable computing architectures that have the necessary flexibility to adapt to different communication environments. Such architectures have to be able to provide the necessary compute power, while satisfying constraints such as cost and power.

Power consumption problems are the result of clock speed to achieve desired bandwidths and frequencies. The architectural dilemma that confronts designers today is that processor speeds are increasing steadily, but memory speeds and access times are lagging. The result is that a simple increase in processor speed even of a factor of tenfold or more may not result in a significant increase in system performance. The solution to this fundamental problem determines tomorrow's designs.

The use of caches: primary, secondary, streaming, etc. is the primary way that industry addressed such problems today. Such caches are composed of expensive "fast" memory and located close to the processor. But the ability of these approaches to offer continued performance improvements is severely limited.

We have been looking for other means of breaking through this constantly narrowing bottleneck. This was the focus of much of the work in architecture at Stanford.

Alternative to conventional processor approaches is the creation of highly scalar, deeply pipelined machines (Super-Scalars) and Very Long Instruction Word (VLIW) machines. At Stanford and elsewhere both of these approaches are pursued, recent industrial machines by Intel (Pentium) and TI (TMS) are of the VLIW type. The real issues involve the question of how to mix and match these architectural approaches to specific sets of applications, e.g. wireless networks. Intel already announced a new effort to develop processor architectures targeted at (non-wireless) networks.

The current DSP processor approach faces a narrowing memory bottleneck, resulting in a high power consumption, limiting the range of applications to ones with relatively low frequency signals. Improvement in power consumption and processor speed may appear as part of the normal evolution of chip technology. However, system performance and system power consumption will still be limited by the memory bottleneck.

In deeply pipelined machines, processing elements are arranged in such a fashion that serial processing occurs in serially arranged processing components. The output of each component is directly connected to the input of the next component. Thus there is no memory latency when moving from one processing step to another. Systems do not have to run at clock rates that are high multiples of the data rate, because each element only has to process data at the data rate. Each sample is processed essentially in one sample time. This reduces power consumption and raises the ceiling on the frequency of signals that can be processed.

Security

The overriding difference between DOD and non-DOD requirements lies in the security area. In non-DOD areas cost-benefit drives the investment away from resources for security. In DOD applications user and mission requirements dictate the minimal level of security. Ideal architectures for such applications require a proper mix of standard and re-configurable subsystems, components, and supporting software. In re-configurable architectures, once unnecessary configuration data is removed it is impossible to determine what the system was doing. Therefore it is very difficult to reverse engineer a reconfigurable system where the hardware is not security algorithm specific.

Such applications require new approaches to data security and protection. These approaches are partly necessitated by new technologies that support wireless networks. Data security and protection methods commonly used in today's wireless networks are likely to be inadequate in future wireless networks. For instance, the increasing number of different services and higher data rates make reliable data transfer protocols that rely on retransmissions become unusable when the amount of data that must be buffered becomes excessive. Therefore data protection in data compression in wireless networks will depend largely on forward error correction and coding methods. These techniques must adapt to a variety of data loss scenarios, ranging from relatively low random bit error rates to total loss of large portions of a transmission. The different services and higher data rates constrain the choice of error protection methods to those of small circuit complexity, and low power.

The different services and variable data rates also require new approaches to data security. Encrypted data must be protected against errors using a higher layer of error correction or error detection and transmission. Data security schemes that can be unified with data protection are appropriate in this new environment.

Beyond data link security and reliability, communicating large data messages in wireless networks based systems is a key performance, functional, and cost issue.

Adding intelligence to wireless communication systems is required. If one can re-configure the mobile units of a wireless data communication system, one can provide significantly improved functionality. One can support major enhancements by making such systems dynamically re-configurable, while keeping the power requirements low:

- Enhance security of data transfer
- Insure integrity of data transfers
- Carry out compression and expansion of data, and multi-media.
- Space-Time Code Diversity, especially over multiple services.

This processing can be done simultaneously with the basic functions required for wireless network services, as well as with higher-level data operations (e.g. multi-media, data fusion) at the source, possible relay, and destination.

Future very-high bandwidth wireless networks will most likely involve new technologies, such as GaAs, photonic, quantum devices, low-power. Early involvement in these technologies will enable compatibility, early insertion, and other synergies that are invaluable in their own right, aside from being "cost shared".

Algorithms for Coding and Security

The characteristics of wireless communication networks present both challenges and opportunities in the area of data-reliability and security. Challenges arise from the highly variable data rates and latencies. Simple protocols such as ARQ (Automatic Repeat reQuest) require large buffers to store real-time data for possible retransmission. Opportunities are provided by the ability to perform dynamically reconfigurable functions and the communications redundancy supplied by variable data rates and multiple independent transmission paths.

The primary reliability and security needs to be addressed are:
- Error detection at high data rates
- Forward error correction to reduce retransmission requirements
- Data security to deal with misrouting and interception

The investigation of integrated approaches that solve all three data protection issues using a unified coding method is required.

In principle, reliable and secure communication of data can be accomplished at low data rates by standard error protection coding and cryptographic techniques applied to the 1D bit streams. In practice, we anticipate that more effective methods will take into account the variable nature of the wireless transmission channel and the digital data.

Data security and reliability can be enhanced using multiple independent data paths. If data is transmitted over a number of different routes so that the intended receiver has a larger probability of receiving each transmission than an eavesdropper, then secret-sharing and other space-time code diversity techniques can be employed to obtain reliability and security.

The simplest secret-sharing method is the following. The information sequence X is transmitted as two independent sequences Y and Z, where Z consists of random or sufficiently complex pseudo-random bits generated by the transmitter and $Y = X + Z$ is the bitwise exclusive-or of the data with this noise sequence. The receiver can easily reconstruct the transmitted data from the two received streams, since $X = (X + Z) + Z = Y + Z$.

Because the random sequence Z is independent of the data X, the sequences Y and Z separately provide no information about the data, so an adversary that intercepts only one of

the sequences learns nothing about the data. (Note that the receiver can perform the reconstruction in real time if Y and Z are sent simultaneously and if the receiver can adjust small relative delays in the incoming data.)

This simple example breaks the secret X into two pieces, Y and Z, both of which are needed to reconstruct the secret, and either of which gives absolutely no information about X. (In fact, Y is a one-time-pad encoding of X, where the key Z is independent of X.) Hence it is a 2-out-of-2 secret sharing system, sometimes denoted as a (2,2) secret sharing system. In general, an (s, n) secret sharing system breaks the secret into n pieces, any s of which allow perfect reconstruction of the secret, and any $s-1$ or less provide no information about the secret. While there are some limits on the allowable values of s and n, for all practical purposes, secret sharing systems exist for all values of s and n, provided of course that $s <= n$.

Communication systems with a number of possible paths between sender and receiver (e.g., the Internet, the proposed National Information Infrastructure, and wireless networks) offer an interesting, non-cryptographic security possibility based on secret sharing. The data to be protected during communication is secret shared in an (s, n) system. If s or more of the pieces are successfully communicated to the receiver, it can reconstruct the data perfectly. But an adversary who can eavesdrop on only some of the data paths learns absolutely nothing about the data, if he intercepts fewer than s pieces.

This simple approach suffers in that each piece is as large as the secret itself. Hence, if a 10-out-of-20 secret sharing system is used with 1 kbit of data, 20 kbits must be communicated, and 10 kbits must be successfully received. Other variations on secret sharing may be more promising for DOD applications and deserves more study. Again using 1 kbit of data, X, as an example, the transmitter generates 20 linear projections of the data $P_1, P_2, ... , P_{20}$, each *100* bits long, where $P_i = A_i X$ and A_i is a *100* times *1,000* binary matrix. For the moment, think of the A_i as being chosen at random, but then being publicly disclosed. The A_i can be thought of as public information used by any two parties wishing to communicate securely. Hence the A_i cannot be kept secret from an adversary, who may, at another point in time, be an ally. While this model is closer to commercial than military applications, even in the military there are restrictions based on need-to-know that bring this model into play.

Classical information theoretic arguments show that any *10* of the P_i allow almost perfect reconstruction of the 1 kbit of data, X. This is because a randomly chosen $n \times n$ binary matrix A has approximately a *29%* chance (*1/2 x 3/4 x 7/8 x 15/16 x ...*) of being invertible and with very high probability A has small rank defect (e.g., the probability of the rank defect being greater than *10* is less than *0.001*). Modulo this small rank defect possibility, the receiver can reconstruct the data provided it successfully receives at least 10 of the 20 pieces. While the small rank defect problem needs to be dealt with, for simplicity of exposition in this proposal, we can ignore it. It appears that careful selection of the A_i can avoid the problem.

Secret sharing by linear projections can be used to protect data that is transmitted over multiple paths. Error correction can be accomplished when more than the minimum number of data pieces is received. However, in practice, random choices of the A_i matrices will not lead to computationally feasible or efficient reconstruction procedures. Therefore investigations into simplified secret sharing and other space-time code diversity methods are needed.

2. Proposed Work

We proposed to carry out research in *algorithms and reconfigurable architectures* for the Adaptive Computing Systems (ACS) program, with *applications to wireless networks,* including *coding and security.* In addition, we proposed research and development of a *reconfigurable hardware test bed* aimed at deployed mobile wireless services, under the Global Mobile Information Systems (GloMo) program.

We proposed research into Algorithms for Secure, Robust, Smart, and Dynamic Re-Configurable Wireless Networks. In order to cope with the evolution and complexity of user and mission requirements, flexible architectures with intelligent and reconfigurable networks and computer systems and subsystems are required.

New coding algorithms and coding techniques for digital data were investigated as part of a framework that provides the necessary network and computer reliability and security for digital data transfers, over wireless communication links. Our framework provides the flexibility, consistency, and independence of assuring reliability and security in various domains and at various levels of the network and computer systems.

Wireless communications present both challenges and opportunities. Our new algorithms take significant advantage of these opportunities. For instance, a unique set of new data protection algorithms is based on secret sharing. These algorithms exploit the available parallelism in the flexible combination of space-, time-, code- and frequency-division multiplexing of wireless network and computing architectures. This leads to an integrated approach that includes the solution to three data protection issues: error detection, forward error correction to reduce retransmission requirements, and data security to deal with misrouting and interception. The ability to trade off coding, network, reconfiguration, and other system parameters enables flexibility, and cost effectiveness.

Security and robustness is supported among other aspects of our architectures, by flexible, table-directed methods at the hardware level via *dynamic re-configuration,* e.g. FPGAs, software reconfigurable gate arrays. Security is enhanced since the hardware is neither algorithm specific nor instruction set dependent. Robustness and security are enhanced by the fine-grained nature of the primitive functions provided by the hardware. They can be used to enhance the security of data transfers, insure the integrity of data transfers, and carry out compression, decompression, and other functions such as future services, e.g. multi-media or communication type data.

These fine-grained primitives are ideally matched to support wireless communication (encoding, transmission, receiving, decoding, security, and future functions). The primitive functions not only directly implement lower level functions required by most data security algorithms, but also directly support higher order functions that improve efficiency and performance of algorithms and associated compilers.

Our study of well matched *re-configurable architectures* and proper levels of granularity of macro functions will help determine future generation architectures of high-performance and low-power commodity FPGAs (instead of being determined by "yesterday's customers typical glue-logic"!)

Our proposed work was also synergistic with ongoing work on security under BMDO/ARO, and other government agencies, for a global high-speed photonic back-bone and satellite networks. The appearance of photonic technology for RF links ("fiber-radio", radar, lidar, image, and other sensors), is not only be compatible and complementary in the short range, but even a crucial component technology of choice for future very high-bandwidth wireless networks. Photonics will play a crucial role in (parametric) RF antenna arrays. We were and continue to be involved in photonic and quantum devices research under Darpa (Ultra), BMDO/ONR, and ARO (MURI99) in terms of potential systems applications, and we are expanding the synergisms among these efforts, including the potential impact on wireless

network services and architectures. For example, we expect that the photonic and quantum devices work will lead to very effective implementations and combinations of very low power communications, quantum-crypto systems, using space-time code diversity.

Our proposed research efforts in the areas of data *security and reliability,* and intelligent processing for wireless networks are described in the following.

Data reliability and security

We proposed research into data security and reliability will concentrate on development of new algorithms and coding techniques for safe and accurate delivery of digital data from diverse sources over highly variable wireless communications links.

Security and reliability via secret sharing

Previously we had assumed a noiseless channel so that received packets (pieces) were identical to those transmitted, although there was a possibility that a packet could be lost (hence we used of a 10-out-of-20 system). Therefore we considered extending the model to allow errors within received packets. If only 10 packets are received, we have only 1 kbit of information about the 1 kbit of data, and we cannot correct (or even detect) errors. If 11 packets are received, instead, we then have 100 redundant bits and error correction becomes possible.

Because random parity check codes achieve Shannon's channel capacity, the 11 received packets allow reliable transmission as long as the error rate is small enough that the capacity of the channel is somewhat greater than 1 kbit/1.1 kbit = 0.91. If 12 packets are received, they allow reliable transmission so long as the error rate is small enough so that the capacity of the channel is somewhat greater than 1 kbit/1.2 kbit = 0.83 , and so on. Hence, from the Shannon theory point of view (in which computational requirements are not considered), this generalization of secret sharing also allows optimal error correction if more than the minimal number of packets is received.

Unfortunately, correcting errors with an arbitrary parity check code is an NP-complete problem, so the computational requirements of this otherwise optimal system can be horrendous. We therefore propose investigating whether computationally efficient error correcting codes (e.g., BCH and Reed-Solomon codes) can be used for generalized secret sharing, and if not, trying to find new classes of codes that, while less efficient from a purely error correcting point of view, also allow generalized secret sharing.

Related to this topic is the question of the tradeoffs between net data rate and coding complexity. We wish to determine, for a given level of reliability, the coding complexity as a function of data rate, as the data rate ranges from very small to the maximum channel capacity. In particular, it is important to determine whether the coding complexity falls off rapidly or slowly as the data rate is reduced from capacity.

Low communication-complexity error detection codes

In addition to the application of secret sharing to network data security and reliability, we propose to investigate topics related to protection of high speed, high volume data communications.

Whatever higher level data recovery methods are used, data packets must be supplied with a checksum that guarantees detection of all but a small fraction of possible errors. The standard method, cyclic redundancy check (CRC), has many advantages. In particular, it is optimal with respect to burst error detection, and can detect all but a fraction, 2^{-p}, of possible errors, where p is the number of check bits. The two most common shift register circuits for calculating CRCs are quite simple and have either constant or $O(log\ p)$ gate delays. However, these circuits require either large fanout or fanin, so that in very high-speed operation, the

long wires needed to communicate the state of the circuit to all parts of the next state logic are the computational bottleneck. We will study the error detection efficiency of simplified CRCs and other linear block codes in which each check bit depends on only a small number of nearby shift register bits. Among the possibilities are ladder-like structures using X-gates instead of shift registers.

Error detection via modulation

We planed to explore how to use redundancy inherent in various modulation schemes to provide sufficient error detection.

Other techniques that are not yet used in the consumer domain, such as low probability of intercept methods (e.g. spread-spectrum, ultra-wide-band signals, femto-second, and quantum-limited technologies) appear to be compatible with our security algorithms, given our current work with BMDO/ARO and other government agencies

3. Project Organization

The research support for this project came from two Darpa programs: The algorithms and reconfigurable architectures research was supported by the Adaptive Computing Systems (ACS) program, and the research into applications to wireless networks, including coding and security was supported by the Global Mobile Information Systems (GloMo) program.

Stanford University was the prime contractor, responsible for research in algorithms and reconfigurable architectures for Adaptive Computing Systems (ACS), with applications to wireless networks, including coding and security.

enVia was a subcontractor, responsible for research and development of a reconfigurable hardware test bed aimed at deployed mobile wireless services (e.g. AMPS, PCS)

Late into the contract enVia was reorganized, but the reorganized enVia/Morphic was transparent to the research contract, as all work was performed under enVia's subcontract; however, the final report is organized recognizing this change of responsibility.

Work and Funding Timeline:

Work actually started on this project on September 6, 1996.

The subcontract with enVia was finalized in mid December 1996, therefore the radio subcontract with enVia started three months late. This subcontracted work roughly maintained this three month delay, some of the work caught up with the original schedule, and some other work got further delayed. Some of the additional delay was caused by commercial partner (Bell South) funding of enVia that was either delayed or finally did not materialize, hence enVia had to wait for other funding sources.

A BMDO security contract effort sharing that was in place with Stanford University, was rescinded by Congress in October 1996, for non-technical reasons; the compensatory student support was later approved, and later funded in part.

4. Research

The Stanford Architecture and Arithmetic group carried out research into algorithms and reconfigurable architectures for Adaptive Computing Systems (ACS), with applications to wireless networks, including coding and security.

enVia carried out research and development of a reconfigurable hardware test bed aimed at deployed mobile wireless services (e.g. AMPS, PCS.)

4.1 Stanford Research

Algorithms, Reconfigurable Architectures, and Security

Introduction

The Stanford reconfigurable architecture research in ACS is based on our PamBlox approach to FPGA programming, an object oriented design methodology that has shown some significant advantages. Various arithmetic and security algorithms were implemented in efficient low-power FPGA designs, including DES and IDEA encryption standards.

The security work has demonstrated a multi-path approach to secure wireless data transmission in 1998 using three laptop computers, based on wired, modem, IR, or Ethernet links, and commercial wireless Ricochet or Freewave links.

We investigated new space-time code diversity based algorithms for wireless communication in the context of software defined radios, with the goal of providing or extending simultaneous security (privacy, integrity, reliability, availability, and protection from reverse engineering) by exploiting parallelism and scalable hierarchical multiplexing schemes to build software defined radio architectures. We developed new Space-Time-Code Diversity (STCD) based security architectures that support all these requirements and map well onto software defined radio technologies.

4.1.1 FPGA---Adaptive Computing Systems Summary

Our Darpa ACS project specific contributions for the FPGA design environment include a set of tools and methods: PAM-*Blox*, consisting of 2 levels of abstraction (PamBlox and PaModules). Our *PAM-Blox* software v1.0 is available to the DARPA ACS community and can be downloaded from the web at   http://umunhum.stanford.edu/PAM-Blox/   This software includes a number of *PAM-Blox* designs for arithmetic (multipliers, matrix-multiply, Jacobi relaxation), and standard encryption (DES, IDEA), that compare favorably with the RAW benchmarks, see our publications below. Generally, we achieved performance improvements in throughput, area, and/or MOPS/Watt, ranging from a factor of two to six of our *PAM-Blox* based FPGA designs over Micro Processors, DSP, and ASICs, in latency tolerant applications. *PAM-Blox* provides a clean interface to share FPGA designs. Therefore we encourage contributions in the form of PamBlox, PaModules and entire applications using these modules. Contributions should follow the style and format given in the base distribution, and contain at least one sample application where the modules are used. All contributions to *PAM-Blox* have to be distributed with the GNU General Public License. Serious contributions will be reviewed and, if acceptable, included in further releases of *PAM-Blox*.

Programmability, Performance and Power Metric for FPGAs

PAM-Blox consists of open, object-oriented circuit generators on top of the PCI Pamette design environment, PamDC. *PAM-Blox* was first introduced in Mencer [MMF Apr98].

PAM-Blox are intended to be part of an open repository that enables design sharing between members of the adaptive computing community.

High-performance FPGA design for adaptive computing is simplified by using a hierarchy of optimized hardware objects described in C++. Programmability is improved by using object-oriented techniques such as templates, virtual functions, function overloading, and inheritance, in the specific way outlined in this section. As with object-oriented software, the design of the hardware object interface is critical to the usefulness of the system. In fact, more effort was spend on the iterative design of the interface, than was later necessary to design the hardware objects themselves.

In order to bridge the space between the algorithmic representation of an application and the gate level (lookup table level) we add levels of abstraction, starting at the register transfer level (RTL) which is equivalent to the PamDC level. The "big picture" is shown in figure below.
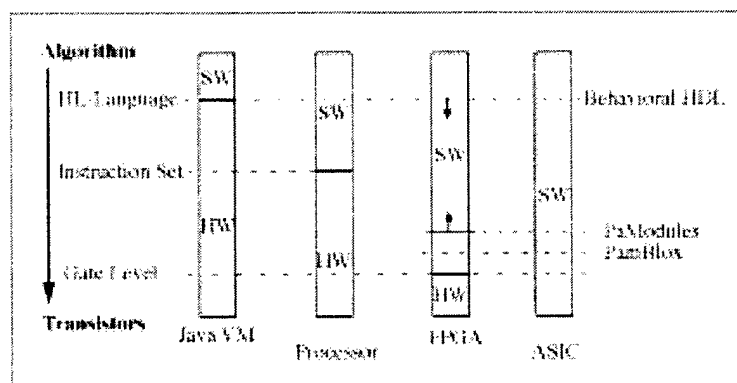
This figure shows the Hardware/Software interface for different technologies: a Java microprocessor (JavaVM), a conventional microprocessor, FPGAs, and ASICs. In general, the programmer has to bridge the gap between algorithms and transistors.

The boundary between hardware and software for processors, FPGAs, and ASICs, defines the interface between programmer/compiler and the computing elements. The low-level boundary between software and hardware for FPGAs requires the software to bridge a large space from algorithm down to the gate level.

The PAM-Story

PAM stands for Programmable Active Memories. The first PAM, PeRLe-0, developed at DEC PRL in France, is one of the first FPGA based computing machines. Next to the hardware efforts the PAM team also developed a C++ class library, PamDC, for creating designs for Xilinx FPGAs. The most impressive result obtained with PamDC and the PeRLe-1 board is RSA encryption at Cray speeds.

The most current PAM is the PCI Pamette board developed by Mark Shand at the DIGITAL Systems Research Center. The PCI Pamette consists of 5 Xilinx XC4000 series FPGAs.

PAM-Blox: PamBlox and PaModules

PAM-Blox consist of two major layers of abstraction. First, *PamBlox* are parameterizable simple elements such as counters and adders. Automatic placement of carry chains and flexible shapes are supported. PaModules are more complex elements possibly instantiating PamBlox. PaModules generally have fixed shapes and are usually optimized for a specific

data-width. Examples for PaModules are multipliers, Coordinate Rotations (CORDICs) [MMF Jun98], special arithmetic units for encryption, and communication modules.

The key difference of our approach to most other design tools for FPGAs is that the designer has *optional* control over placement at each level of the design hierarchy, which is the key to high-performance FPGA design. As mentioned before, the object interface was chosen carefully to encourage code-reuse and simplify code-sharing between designers.

A subset of the PamBlox hierarchy is shown in the figure below
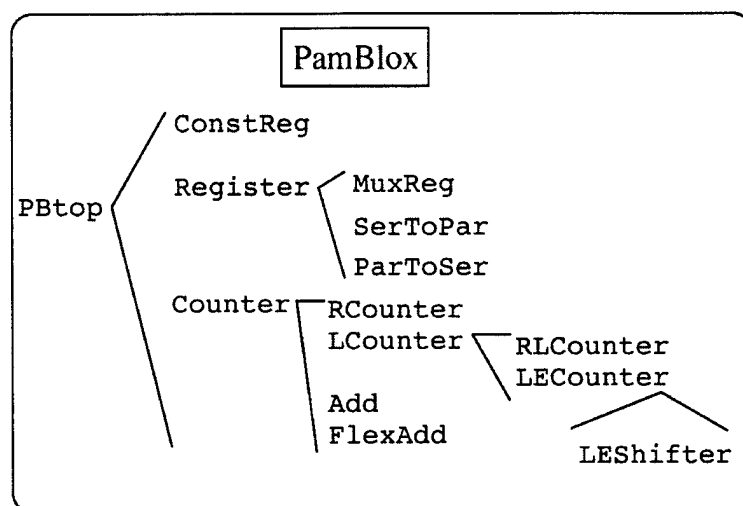
representing a subset of the PamBlox hierarchy. The top object PBtop consists of a minimal amount of logic and a set of placement functions that are inherited by all PamBlox objects. Prefix `R' stands for "Resetable", `L' for "Loadable", and `E' for an "Enable". The top object, *PBtop*, consists of a vector of registers, and a set of placement functions that handle different carry-chain configurations. As an example of code reuse, every child of *PBtop* inherits the placement functions and can overwrite them if necessary.

PaModules are more complex, generally fixed circuits implemented as C++ objects. PaModules can include multiple PamBlox and are optimized for a specific data-width. Examples are constant(K) coefficient multipliers (KCMs), Booth multipliers, Coordinate Rotation Digital Computer (CORDIC) circuits (mentioned above), and special purpose arithmetic units such as a constant multiply modulo $(2^{16} + 1)$ operation for encryption.

The table below shows the code-size of *PAM-Blox* version 1.0 circuit generators. Code-size is given in PamDC / C++ lines necessary to implement the objects.

	PamBlox	PaModules
No. of Objects	28	6
Lines of Code	1370	750
Av. Lines per Object	~ 50	~ 120

Hierarchical Naming

Commercial synthesis tools make it difficult to optimize circuits on a low level (e.g. after place-and-route) by flattening the design and changing the naming of the wires. PamDC

enables direct control over the naming of wires. *PAM-Blox* are implemented to support a hierarchical naming scheme that creates a unique name for each wire in the design similar to paths in a file-system. The name of each wire contains all the ancestors (parent objects) of the wire. The top name can be specified by the designer, e.g. a PaModule multiplier with the name "multy" containing one adder with a carry-chain, results in the following name for the third element of the carry-chain, e.g. *multy/add0/carry<3>* .

The naming scheme enables designers to use additional tools for debugging and still be able to trace the source of each wire found in the final netlist. For example, the naming hierarchy is preserved for simulation (within PamDC) and low-level tools such as Xilinx fplan.

PAM-Blox: Performance

Applications that have been shown to execute favorably on FPGAs are data intensive applications which can be executed in very deep pipelines (e.g. encryption, pattern matching, etc.) and applications with a huge amount of fine grain parallelism such as Jacobi relaxation and lattice gas simulation.

Given today's technology, FPGA based computing machines can compete with general purpose processors on latency tolerant applications that require a relatively small amount of logic during a specific period of time, which we refer to as *persistence* of the associated task. *Persistence* has to be an order of magnitude larger than reconfiguration time. Due to large reconfiguration times of today's devices single FPGAs do not scale easily to large problem sizes or large data-flow graphs. Multiple FPGAs can be used to compute larger problems. The major drawback is the very high complexity of partitioning a design onto multiple FPGAs given a limited amount of pins. Overcoming the pin-limitation in software -- with design tools -- was investigated in the Virtual Wires project at MIT. Eliminating the pin limitation with multi-chip modules of FPGAs was explored in the Teramac project at HP.

First, we compare the original implementations of the RAW benchmarks Jacobi, Matrix Multiply and DES encryption synthesized by Synopsys FPGA Express II with implementations using *PAM-Blox*. The *PAM-Blox* implementations of these RAW benchmarks have been designed by trying to keep the design effort within order of magnitude of the design effort for the behavioral implementations in Verilog. Note that *PAM-Blox* supports optional, hierarchical placement, while the design methodology of FPGA Express does not enable placement by the designer. Ideally, object-oriented hardware generators will be combined with object-oriented behavioral CAD tools.

Our metrics of comparison are *minimal cycle time* and *area requirement* in configurable logic blocks (CLBs), and *compile time* (from C++ / *PAM-Blox* to the Xilinx Netlist Format). Area can be shown to be directly proportional to power consumption, giving us an additional perspective to the data presented below. Compile time was measured on an Intel Pentium PC at 120 MHz. *PAM-Blox* are compiled with Visual C++ 4.0. The objective of this work was to put real results from *PAM-Blox* into perspective with a state-of-the-art, commercial CAD tool -- not to be mistaken for an argument for structural (*PAM-Blox*) versus behavioral (FPGA Express/Verilog) design -- while acknowledging that it took more effort to create the *PAM-Blox* designs.

We present the performance of the following three RAW benchmarks:

A. Jacobi Relaxation

Jacobi relaxation is an iterative method for solving differential equations of the form:
$$-^2A + B = 0$$
The basic operations for this benchmark are shift and add. The implementations compared in the table of RAW benchmarks below consist of a 4x4 array with 2x2 active cells and 8 bit values. During each clock cycle, each active cell takes the values of cells neighboring east, south, west and north, adds them together and divides the result by four. The arithmetic

operations 'shift and add' map easily onto the Xilinx XC4000E library used by FPGA Express II. Therefore there is not much room for area improvement. Clock frequency of the *PAM-Blox* design is only about 15% higher than the design optimized by Synopsys HDL compiler. However, the improvement in area is about 20%.

B. DES Encryption

DES encryption is well suited for implementation in hardware. The basic primitives are fixed permutations and exclusive-or. The results for the *PAM-Blox* DES design in table below show a 30% increase in performance (clock frequency) using only half of the area.

The superior results obtained with *PAM-Blox* are due to partially manual placement and technology mapping, i.e. the careful design of logic that fits into 4 bit lookup tables.

C. Integer Matrix Multiply

The Matrix Multiply benchmark multiplies two 4x4 matrices with $4^2 = 16$ multipliers and an adder tree. FPGA Express II uses simple bit-serial shift-and-add multiplication. A full matrix multiply therefore takes more than 50 clock cycles. For this benchmark we chose to create a more efficient computational structure, i.e. an arithmetic unit to show how *PAM-Blox* can be used to adapt the arithmetic units to the specific requirements of the application. By implementing multiple bit-serial multipliers using Booth encoding, we are able to trade area for performance. Obviously Booth multipliers are more efficient for this specific application. The idea is to use the *PAM-Blox* environment to choose the arithmetic unit -- in our case the multiplier -- that is best suited for the specific application.

RAW Benchmark Results

The Table below shows two PAM-Blox designs, PAM-Blox 1 and PAM-Blox 2, differing only in the selection of the multiplier PAM-Blox 1 multiplies the matrices in 27 clock cycles while PAM-Blox 2 takes 19 clock cycles for a full 4x4 matrix multiplication including data transfer. Clock cycle times for the PAM-Blox designs are around 33 MHz. The original design synthesized with FPGA Express II runs at 15 MHz and requires 39 clock cycles for a full matrix multiply.

|	Compile Time	Area [CLB]	Frequency
JACOBI 4x4 (8 bit)			
FE II	80 s	164	30 MHz
PAM-Blox	45 s	129	35 MHz
DES (1)			
FE II	1,510 s	828	15 MHz
PAM-Blox	86 s	398	22 MHz
MATMULT 4x4 (8 bit)	Compile Time	Area [CLB]	Mega-mmps
FE II	350 s	609	0.38
PAM-Blox 1	77 s	604	1.23
PAM-Blox 2	98 s	954	1.52

RAW benchmarks compiled with Synopsys FPGA Express II (FE II) are compared to *PAM-Blox* implementations. Compile time stands for the time to compile a design description to a Xilinx netlist file. FPGA Express results are reported for the completely placed and routed system. The performance for matrix multiply is given in matrix multiplications per second (mmps).}.

This table shows the throughput in matrix multiplications per second.   With the right multiplier we see an increase in throughput of up to 4 times, compared to the original RAW benchmark compiled with Synopsys FPGA Express II, using a generic shift-and-add multiplier. Of course selecting a faster multiplier yields better results, but PAM-Blox designs can compete with commercial CAD tools.

D. Constant(K) Coefficient Multipliers - KCM

Constant(K) Coefficient Multipliers (KCMs) are of interest for many applications including filters and encryption. KCMs are implemented as PaModules. We compare 16 bit KCMs with a throughput of 16, 4 and 1 bits per clock cycle respectively, in the figure below, showing Area and Latency for 16 bit Constant(K) coefficient multipliers with a throughput of 1,4 and 16 bits per clock cycle. Latency is shown in units of 10^{-1} clock cycles i.e. 160 means 16 clock cycles.

This comparison demonstrates the time-space tradeoff for KCMs on Xilinx XC4000 FPGAs. With increasing throughput, we increase the area requirement and decrease latency -- trading area for latency and throughput. While this is not surprising, the generator framework allows the designer/CAD software to easily choose and modify the arithmetic unit that fits the specific requirements of the design problem.   First, we implemented the fully-parallel KCM with *PAM-Blox*, achieving the same performance and area values reported in the application note from Xilinx. Second, we created a digit-serial design for a 16 bit KCM which takes 4 bits at a time at about *1/3* the area of the fully parallel version. The design of this multiplier is related to distributed arithmetic, combining multiply-adds into table lookups.   While performance over area for this multiplier is worse than for the fully parallel case, the small area of this multiplier allows us to map an entire multiplication-based encryption algorithm onto around 3200 CLBs. The relatively small size of the bit-serial KCM (17 CLBs) allows us to fit more than 45 such multipliers on a Xilinx XC4020E with 800 CLBs.

E. Compile Time and Xilinx Place-and-Route

The improvement in compile time for the RAW benchmarks above is a consequence of interpreting Verilog versus direct execution of C++, and optimizations within FPGA Express II. Except for DES, the benchmarks are simple and require almost no optimization. FE II increases structural-like compile time by about half an order of magnitude.

We expected the manual pre-placement to decrease the remaining automated place-and-route time. Instead we found that Xilinx place-and-route performance is dominated by

routing. Place-and-route performance therefore varies depending on how easy or hard it is to route the placed design.

While hand-placement improved circuit performance, place-and-route times varied depending on the specific design, FPGA size and seed number for the non-deterministic place-and-route algorithm.

4.1.2 Coding and Security

Introduction

We investigated new space-time code diversity based algorithms for digital radios / wireless communication in the context of software defined radios, with the goal of providing or extending simultaneous security (privacy, integrity, reliability, availability, and protection from reverse engineering) by exploiting parallelism and scalable hierarchical multiplexing schemes to build software defined radio architectures. We developed new Space-Time-Code Diversity (STCD) based security architectures that support all these requirements and map well onto software defined radio technologies.

Background

Understanding the interactions of Coding and Security, Software Defined Radio Architecture, Simulations, and Basic Technologies is our basic goal.

A list of the topics studied includes: network security, primary security requirements, standard approaches to security requirements, security via secret sharing, use of space-time-code Diversity, reliability using multiple paths, integrity using multiple paths, privacy using multiple paths, Shamir's polynomial interpolation method, unified reliability, integrity, and privacy, network and routing issues, and ramp schemes.

Multiple Path Security Demonstrations

In February 1998 at Stanford during a site visit by Dr. J. Munoz, we demonstrated, according to plan, a two service, two link type security algorithm (secret sharing) running on laptop computers, see Appendix A. Encryption security algorithms DES and IDEA were implemented with *PAM-Blox* on FPGAs.

We also tested commercially available wireless link hardware (Freewave and Ricochet) for their suitability for testing our security algorithm implementations with RF links. We tested of various TCP/IP and SLIP/PPP point-to-point connections using different hardware and software combinations. We tested the potential for the simultaneous use of two different wireless (Freewave and Ricochet) and two different wired links (Ethernet and modem), i.e., the potential for up to 4 parallel hardware links operating in hardware concurrently, and a number of virtual links per hardware links. The software constraints are given by the standard TCP/IP and OS (NT, Linux, MacOS) limits.

For example, we combined standard protocols such as FTP and PPP (e.g. using FreePPP with Ricochet), and hardware links such as direct serial, serial modem lines, serial wireless connections (Freewave and Ricochet), Ethernet (TCP/IP on Unix, PC, Mac, and AppleTalk among Mac's. We established a number of multiple simultaneous hybrid links between different type of machines and OS'. Numerous experiments were carried out besides using our PC laptop demonstration hardware. For instance, we operated 4 simultaneous hybrid hardware links on an older PPC Mac supporting 11 simultaneous logical links (file transfer, web browsers, and telnet) to Unix fileservers, and PC laptops running Timbuktu on PC's and Mac's. Downloading up to 8 files in parallel achieved close to maximal throughput on wireless (Ricochet) links, 48kb/s out of 56kb/s max. PC's were able to achieve over 70kb/s FTP throughput out of 115kb/s max (ISDN rate) on wireless (Freewave) links. Mac serial link throughput was generally limited to about 50% of max rates due to OS overhead.

The systems integration effort required to test our security algorithms dominates the effort to implement particular security algorithms. We were studying a library approach to deal with security algorithms and system integration.

A number of technical reports and papers (e.g., for conferences and workshops) were generated, see the list of publications below. We pursued a number of efforts to cooperate with other members of the DARPA ACS program.

Multiple Path Security Studies, Summary

The objective of this work was to study the interactions of Coding and Security, and Basic Technologies in the context of software defined radios.

Coding and Security

Coding methods that are appropriate for data security in wireless networks were investigated. These networks have several characteristics that distinguish them from other currently employed networks. The most significant characteristics are unreliable and variable bandwidth channels; predominance of voice, with data and video to follow; to come: narrowcasting -- transmission of data from one source to a designated set of receivers, and data fusion -- combining related data from several sources; simple sensor nodes with limited buffering. These characteristics affect both the lower level network design and the higher level coding methods. Data security encompasses *privacy*, *integrity*, *reliability*, and *availability*; *robustness* from reverse engineering or node capturing is an additional set of issues. Privacy, integrity, and reliability can be provided through encryption and coding for error detection and correction. Availability is primarily a network issue; network nodes must be protected against failure or routed around in the case of failure. A number of additional security issues can benefit from network redundancies. One of the more promising techniques is the use of *"secret sharing."* We consider this method as a special case of our new space-time code diversity (STCD) based algorithms for secure communication. These algorithms enable us to exploit parallelism and scalable multiplexing schemes to build robust and secure wireless network architectures.

Introduction

Robust and secure wireless networks for voice, data and video communication s impose major challenges for the design of such networks. This involves specifying flexible architectures with intelligent and reconfigurable communication processors to cope with extremely variable link characteristics, and determining coding techniques to provide for security needs, such as privacy, integrity, availability, and reliability.

Wireless networks present many challenges in the area of security. Challenges arise from the variable channel characteristics. The coding and encryption methods that were designed to perform in environments appropriate for electronic communications may be inappropriate for wireless channels. High dynamics in fading or other link statistics make coding for average behavior very inefficient or plain unacceptable. Primary security requirements include:

     _   Reliability---data communications protocols should be robust enough to withstand link and node failures and misrouting;

     _   Integrity---data is not modified accidentally or deliberately tampered with, by replacement, insertion, or deletion;

     _   Privacy---confidentiality of the data in the network should be maintained even if an eavesdropper can tap one or more links in the network.

performance of the processor. VLIW processors are not adaptable to unpredictable or variable system behavior, which is especially a problem when dealing with non-uniform memory architectures. Considering the characteristics of both super scalar and VLIW processors, it is clear that an ideal processor would have the advantages of both and the liabilities of neither.

This research has developed an architectural model as well as a simulation tool that is being used to evaluate performance variations across a design space that spans both super-scalar and VLIW architectures. By varying parameters to both the compiler and the simulator, many different processor configurations can be simulated and their performance compared. The basis for this model is the *Split-Issue* execution paradigm which separates the architectural (or *virtual*) behavior of an given machine (behavior that is exposed to the user) and the implementation (or *real* behavior (behavior that is native to the particular implementation but is not exposed to the user).

An Introduction to Split-Issue

The execution of an operation typically consists of three distinct phases: the acquisition of source values, the computation of result values, and the delivery of the completed results. In a traditional pipelined processor, these are often described as *register fetch*, *execute*, and *register write-back*, and are considered to be inseparable steps in the execution of an operation in the pipeline. In the Split-Issue model, these phases are decoupled from each other so that they can be treated independently in the implementation. In order to maintain the correctness of the virtual machine, each event is scheduled to take effect at the appropriate time based on the architectural specification for each operation. Whereas in the traditional pipeline, source, result, and intermediate values are held in latches in pipeline stages until conditions are set to proceed, in the Split-Issue model these are held in temporary storage locations.

The original concept for Split-Issue was presented by Rau at HP and described as a mechanism for supporting a dynamic execution model for VLIW processor capabilities. In the Split-Issue model, described in detail in [R 94] and [R 95], operations are split into three independent phases: *PhaseOne* performs all acquisition events, *ExecutePhase* performs all computation, and *PhaseTwo* performs all delivery events. The PhaseOne and PhaseTwo events are similar in that they are access events performing read and write accesses to the register file and are synchronized to maintain the virtual ordering of operations as scheduled by the compiler. The ExecutePhase events are computations that operate on the operands that have been fetched by the PhaseOne events and produce results that will be written back by the PhaseTwo events. Unlike the PhaseOne and PhaseTwo events, the ExecutePhase events operate asynchronously whenever their source values are available---all synchronization is maintained by the other two phases. A given operation can be thought of as a triplet of events that will be processed by each of the three queues. Note that the ExecutePhase events do not need any timing information specified, since they are able to commence once all operands are flagged as available.

Arbitrarily complex operations can be constructed in this manner and even the most complex CISC processor operation can be framed in this form. One tangential use of Split-Issue techniques might be the emulation of CISC processors (including high performance DSP and Network processors). Emulating a CISC processor by mapping the CISC operations (the architectural operations) into native operations has recently come into vogue in two commercial microprocessors---both the AMD K5 and the Intel Pentium Pro exploit the use of distinct virtual and implementation architectures. In fact, simply by changing the mapping between the virtual and implementation architectures, it is possible to have a single implementation that can emulate different virtual processors.

An observation of the representation used in the above two examples is that it bears a striking resemblance to the information contained in a machine specification for the operations. In a machine-readable form, this is essential information that a compiler would

decreases linearly down to 0 with each additional compromised share.

This allows for $m = k - j$ secrets to be communicated using n shares, improving the communications rate. In this case we trade improved communication rate/efficiency with reduced security; however, it is important to note that the information obtained by the shares in a ramp scheme is on the set of secrets as a whole and not on any individual secret. If each secret is further encrypted, an eavesdropper will have great difficulty in taking advantage of any partial information obtained.

In summary, our Space-Time-Code Diversity secure communication has several advantages. It

a) provides a mechanism for enabling a three-way trade-off between the aggregate bandwidth, computational overhead, and the level of security;

b) offers a means of communicating with very high levels of security;

c) is a "distributed solution" to security problems;

d) has low complexity implementations well matched to software defined radio networks.

Communication over multiple paths, at different times, and with different coding methods is a unified approach to security requirements. For details see our publications, e.g. [CGMF Nov97].

4.1.3 VLIW Approaches to Reconfigurable DSP and Network Processors

General purpose high-performance processor design has recently taken two different (and often diametrically opposed) approaches. One approach used to design high-performance processors is to increase the execution rate by increasing the clock rate of the processor or by reducing the latency of operations. Another approach is to issue and execute multiple operations concurrently. Traditional processor designs that issue and execute at most one operation per cycle have been referred to as "scalar" designs. Processor designs that can issue and execute more than one operation per cycle will be referred to as *super-scalar* processors.

Reconfigurable processors can be viewed as special cases of micro- or nano-programmable (CISC) processors. FPGAs are very fine grain (gate-level) nano-programmable processors. Some approaches to FPGA programming involve emulating controller, DSP, or RISC core instruction sets. An alternative approach is pursued for instance by the Triscend Corporation by offering FPGAs combined with controller or ARM processor cores. Our former Ph.D. student, F. F. Lee, developed this architecture.

Historically, two primary techniques have been used to achieve super-scalar performance. The first technique uses dynamic analysis of the instruction stream during execution to determine those operations that are independent (processors that use this technique are commonly referred to as super-scalar processors). These independent operations can be issued and executed concurrently. The second technique uses static analysis performed by the compiler to schedule independent operations together into compound multi-operation instructions (processors that use this technique are commonly referred to as VLIW processors). All operations in a given instruction can be issued and executed concurrently with no dynamic analysis.

Neither super-scalar nor VLIW architectures are entirely satisfactory. Super-scalar processors require complex control logic to determine dynamically the operations that are issuable, which limits the amount of parallelism that can be exploited as well as the

There are a number of standard approaches to deal with security requirements; we pursued alternative methods to achieve such requirements that are better matched to software defined radio technologies.

Secret sharing can be used in a unified approach to satisfy these security requirements. Secret sharing takes advantage of the possibility of transmitting shares of a message over different paths, at different times, and perhaps using diverse coding schemes. We call this capability of multiple transmissions {space-time-code diversity}.

Reliability can be achieved using multiple paths, i.e., a message may be transmitted more than once, either in time or space. For example, two copies of messages can be sent over two different paths (e.g., one direct link and one path through an intermediary). Communication is successful if either copy of the message arrives. Security is limited to link-level encryption. If any link is compromised, the message will be intercepted.

Integrity can be achieved using multiple paths, i.e., if a message is transmitted over three different paths, an altered copy can be detected and corrected using the two correct copies. If $n = 2k + 1$ distinct paths are used, we can protect against coordinated tampering with k message copies. (Achieving integrity using such a repetition code for error correction is inefficient.)

Privacy can be achieved using multiple paths, i.e., assume two disjoint paths from transmitter X to receiver Y. If an eavesdropper can tap only one link, then one bit of information m can be sent from X to Y with perfect security. X generates a random bit r and sends the "shares" m XOR r and r over the two disjoint paths. Y reconstructs the message m from the two shares XOR-ing the two received shares, thereby canceling r, i.e., getting back the message m. Neither shares in transit gives any information about m.

Blakley (1979) and Shamir (1979) introduced secret sharing schemes as a solution to the preceding key management problem. The key (the "secret") is broken into n pieces called "shares" in such a way that a) any subset of k or more shares can recover the secret; b) no subset of k - 1 keys contains any information about the secret. Such a secret sharing method is called a $(k; n)$ threshold scheme.

Secret sharing schemes can function in "courier mode". The messages themselves, rather than the cryptographic keys, are broken into shares that are given to different couriers -- that is, they are transmitted over diverse paths in space or time. Secret sharing can provide integrity and reliability in addition to privacy. If we can guarantee that a) at most l shares are lost or delivered with detectable errors, and b) at most t shares are tampered with or arrive with not obviously detectable errors, then a $(k; n)$-threshold scheme can be used to achieve perfect security (privacy, integrity, reliability) as long as $n - k >= 2t + l$. For example, 6 disjoint paths can resist tampering with one share, loss of one share, and interception of two shares.

For a sufficiently rich network, there will be many disjoint paths from transmitter X to receiver Y. New routing algorithms will be needed. Instead of simply finding a good route from X to Y, we must now find n reasonably good paths that do not intersect. The number of logical paths may be increased by use of transmissions at different times. This will improve reliability in the presence of transient failures, but does not in general improve privacy or integrity if a link or node has been compromised.

Practical implementations will most likely take advantage of "ramp schemes". The bandwidth expansion required by secret sharing can be reduced if the requirement that $k - 1$ shares provide absolutely no information about the secret s is relaxed. If instead we use a modified scheme:

a) Given up to j shares, no information is gained.

b) From j + 1 to k the remaining uncertainty about the secret vector

need to schedule the operation and is also essential information that a processor designer would require to implement the same operation. This natural connection between specification and model is fortuitous---and the next section demonstrates that there is a similar connection between the model and implementation as well.

The NV Simulator and Experiment Environment

The simulation environment used in this research is the {NV} simulator, which is part of the University of Illinois (Urbana--Champaign) IMPACT compiler suite. It is based on the Split-Issue model and is parameterized to allow the simulation of a wide range of system configurations and models both the processor and memory systems. The table of simulator parameters lists a number of the areas that are parameterized in the simulator. In addition to the capabilities listed in the table, the compiler is parameterized and supports any configuration of virtual function units and schedules code based on the characteristics of the available function units.

Simulation Parameter Summary

The processor architecture being modeled is that of a fixed-width VLIW processor -- this is the architectural processor. The actual configurations that are used in the experiment vary from a traditional static VLIW processor to a complex out-of-order processor that can have many operations outstanding and dynamic scheduling of available operations.

We are in the process of running a number of SPEC92 and Unix benchmarks under the simulator, varying a number of the parameters to understand better the performance implications changing different aspects of the model. The current experiments consist of varying the number of function units, the number of scheduling structures, and the main memory latency to see how the performance varies over these changes. The initial results have shown that, as expected, as the number of function units per scheduling structure increases the utilization of the function units improves. They have also shown that, again as expected, the number of outstanding operations allowed also increases the utilization of the function units. Both of these situations result in the increase of the operation pool available to a given function unit, and thus improve the possibility that an operation will be available to issue to that function unit.

Future work in this area includes addressing both compiler as well as architectural issues. The current architectural model supports unconditional in-line code execution only. This is a significant limitation on performance and future work will include the ability to have the compiler schedule operations that are predicated and non-excepting for static speculation as well as having the processor follow predicted branch paths for dynamic speculation.

4.1.4 Improvements in Video Compression for Multi-Media Applications

Our research in this area primarily focused on novel algorithms for video compression and developed more effective methods to transmit moving images across a communication system with very limited computational resources. The results of our research [Y Dec99] are expected to have an impact on a number of communications applications:

1. Transmission of images across wireless networks for web browsing or email.
2. Video-telephony, such as videophones or videoconference calls.

We considered three key aspects of a video compression system:

1. Video fidelity metrics and their effectiveness at predicting the visual fidelity of different video system outputs. Such fidelity metrics are important when different types of video systems are evaluated, avoiding the use of extensive psychological viewing. Effective video fidelity metrics are consistent, efficient, and have low computational cost.

2. The video bit rate, proportional to the compression ratio for a given video frame rate, refers to the amount of image data needed to represent a certain amount of information. The bit rate should ideally be as low as possible, while preserving the fidelity. Compact image data representations reduce necessary resource requirements, enabling faster transmission and lower storage.

3. The speed of execution for a video compression algorithm is determined by its computational complexity. An algorithm should carry out a given task within certain latency restrictions, while still computing the same, or virtually the same, result. A robust algorithm tries to minimize both bit rate and computational requirements without compromising video fidelity significantly. We addressed all three aspects of the video compression system in more detail, see [YLF Feb97].

In this research, we assessed the capability of four metrics (average MSE, average SNR, ANSI parameters, and ITS metric) to determine the fidelity of video sequences. To do this, we first defined the ideal requirements for a video fidelity metric in terms of monotonicity, degree of change, and consistent behavior. Then, we constructed a series of highly reproducible degraded sequences containing artifacts common to DCT-based transform coders, such as H.263, and evaluated the performance of each metric on those sequences. From the resulting data, we determined the accuracy and reliability of each of these metrics. Our analysis concluded that the better-established average MSE and average SNR metrics exhibited superior performance over the ANSI parameters and the ITS metric, see [YLF Sep97].

We used a standardized video compression system, known as the H.263 standard, which has optimal performance at low bit rates for video-telephony applications, to develop a more efficient algorithm for video encoding. To do this, we characterized H.263 bit-streams and observed that, often, data units, known as "inter-macroblocks," are reduced to zeros after a two-step process within the system called Discrete Cosine Transform (DCT) and Quantization. We took advantage of this observation and proposed a new H.263-compatible algorithm that predictes when there will be all zeros in inter-macroblocks at that stage. This eliminated a significant portion of computation usually needed for those inter-macroblocks and reduces the requirements for wireless communications. Simulation results show that, relative to the base (original) H.263 encoder, our early-detection mechanism can reduce the computational requirements of the DCT/Quantization by as much as 81 while simultaneously reducing the bit rate, with virtually no visible change in video fidelity, see [YLF Nov97].

We extended the work done previously to improve the H.263 system even further. By optimizing the software used to generate H.263 bit-streams, we reduced the computation time of the encoder by four times its original value. We then used a zero-detection algorithm, similar to the one used for the DCT/Quantization, to improve the performance of another section of the H.263 encoder, known as the motion estimator. Our simulations demonstrate that the combined improvements to the H.263 encoder ultimately reduce the computational time of the original encoder by four to nine times its original value.

4.2 enVia/Morphics Research and Development

enVia was a subcontractor, responsible for research and development of a reconfigurable hardware test bed aimed at deployed mobile wireless services (e.g. AMPS, PCS)

Late into the contract enVia was reorganized, but the reorganized enVia/Morphic was transparent to the research contract, as all work was performed under enVia's subcontract; however, the final report is organized recognizing this change of responsibility.

4.2.1 Morphics Final Report

**Re-configurable Multi-mode, Multi-band Information Transfer Systems
Baseband Multiple-Standard Development Section**

Morphics final report is contained in the Appendix B.

Summary

The Morphics part of the project's initial intent was to show that multiple digital standards can be implemented in a single baseband architecture without multiple, parallel processing paths. This development would then be incorporated with a separate multi-band RF development to create a single multi-standard, multi-band communications device

The scope of the project as originally conceived proved too big to attempt in its entirety. An end-to-end mobile handset is by its nature a $50M project. An attempt to use commercially available platforms was partially implemented in that the development hardware was obtained. However, no software was available, so a complete multi-standard baseband implementation was not attempted.

The *target development* then became a multi-standard hardware logic section, the only new, previously undeveloped part of a multi-standard communications baseband, to enable a *multi-standard digital modem*.

The initial candidate services were AMPS, GSM, IS136, and IS95, i.e. analog cellular, European TDMA, North American TDMA, and CDMA.

The services considered during the scope of the project were all of the TDMA based services: GSM(800, 1800, 1900); IS-54 & IS-136; PDC.

Introduction

The initial intent of this development is to show that multiple digital standards can be implemented in a single baseband architecture without multiple, parallel processing paths. This development would then be incorporated with a separate multi-band RF development to create a single multi-standard, multi-band communications device

The scope of the project as originally conceived proved too big to attempt in its entirety. An end-to-end mobile handset is by its nature a $50M project. An attempt to use commercially available platforms was partially implemented in that the development hardware was obtained. However, no software was available, so a complete multistandard baseband implementation was not attempted.

The target development then became a multistandard hardware logic section, the only new, previously undeveloped part of a multi-standard communications baseband.

Below we highlight parts of the Morphics presentation style final report, see Appendix.

Development Partitioning:

Proposed development of the communications system:

- Third parties to supply the DSP development environment for digital development
- Morphics DRL enables multi-standard digital modems
- enVia enable multi-band RF communications

RF

IF

enVia Technology

Conversion

Multi-standard
Processing Engine

Morphics Core Technology Developments

DSP

Processor

Commercial
DSP Development
platforms

Functional Mapping:

The initial candidate services were AMPS, GSM, IS136, and IS95. These are Analog Cellular, European TDMA, North American TDMA, and CDMA.

The services considered during the scope of the project were all of the TDMA based services:

GSM(800, 1800, 1900); IS-54 & IS-136; PDC

The GSM standards were implemented and the IS-54/136 were in process at time of conclusion of the project.

The CDMA services would have been the next logical development under an extension to the existing program.

Product Content:

Digital Communications: Analog RF, Analog Baseband, Mixed-Signal, DSP (Digital Signal Processors), ASSP (Application Specific Signal Processors), Microprocessors/Micro-controllers.

Balancing Hardware & Software:

The challenge is to efficiently manage the system complexity by balancing HW/SW, i.e. system-level design involves HW/SW module creation & verification.

Mapping System Specifications to Hardware & Software:

The diagram below represents the methodology used to carry out the mapping of system specifications to HW/SW. System specifications are derived from applications and standards; after algorithm exploration, system level partitioning is carried out into hardware (e.g. standard cells, FPGAs, and application specific hard-wired processors) and software (e.g. general purpose programmable processor.)
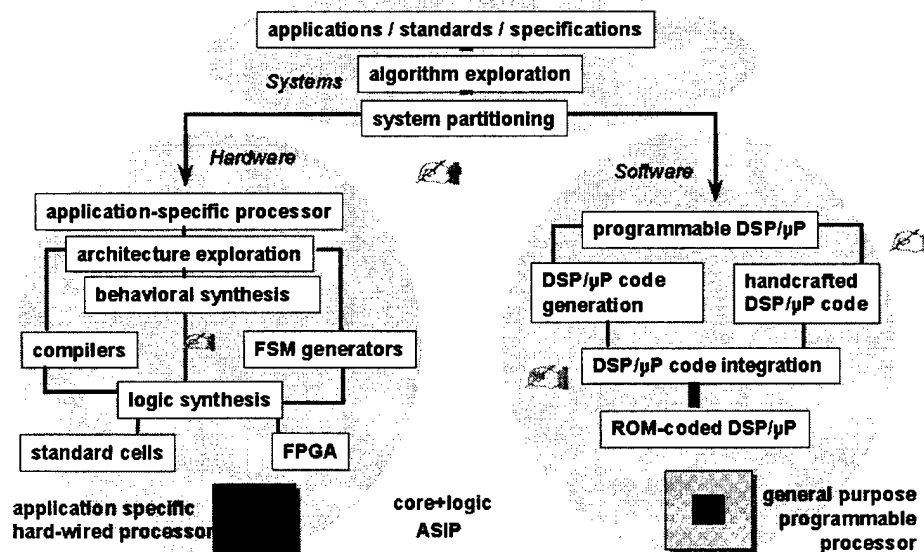
High-Level Multi-Standard Architecture:

Current product architectures in common use support single standards in multiple components, and leave 3 key issues to address:
1. Multi-Standard Baseband Modem.
2. Multiple-Bandwidth Signal Conversion.
3. Multiple-Frequency Radio Transceiver.

Future architectures must support multiple standards, while significantly reducing the number of components; two are viable in the near term, the eventual goal is one.

1. Multi-Standard Baseband Modem

A complete DSP and controller with Modem functionality for multiple standards

2. Multiple-Bandwidth Signal Conversion

An ADC/DAC element that is variable in BW and resolution allowing support of multiple standards.

3. Multiple Frequency Radio Transceiver

Radio Frequency (RF) section capable of operating over the 800/900/1800/1900/2000 bands.
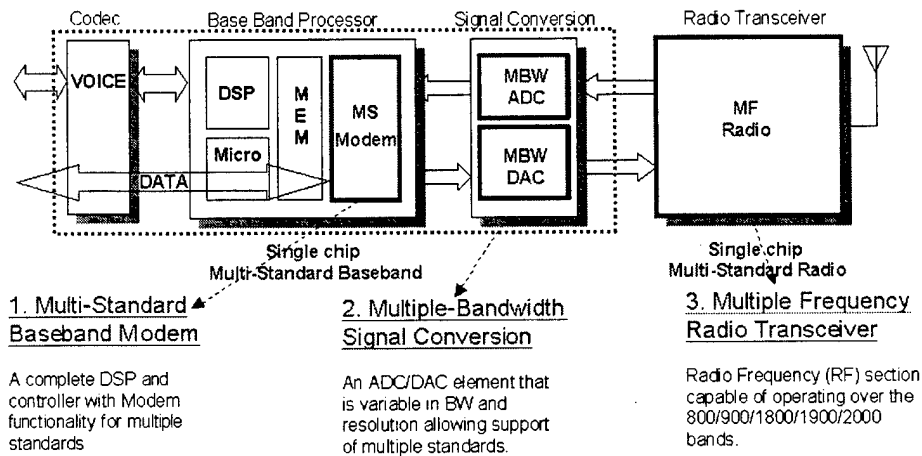
Reconfigurable Architectures:

When the various services are examined, common elements emerge. These lend themselves to implementing in "hard" ASIC cores, with reconfigurable logic for that which is specific to each service

Application Specific Reconfigurable Architectures, "Beyond FPGAs":

Morphics is spending a significant effort into improving the efficiencies of designs.

Specifying a class or the range of applications can significantly improve the required area, power, and performance of reconfigurable architectures *beyond FPGAs* , by matching the flexibility of the reconfigurable architecture to the class or range of target applications. Combining the optimization of macro- and micro-architecture with a suitable design

methodology it is possible to *reduce the required silicon area* by a *factor* of *45* ! A concurrent reduction in power is evident. *This is one of the major results obtained in this contract.*

Flexible Wireless Systems Architectures and the Technology Confluence:

Flexible wireless systems architectures are enabled by a proper combination of DSP-based Soft-Modems, Reconfigurable Logic, and System-On-Chip VLSI Design. This research has been uniquely successful in combining diverse capabilities.
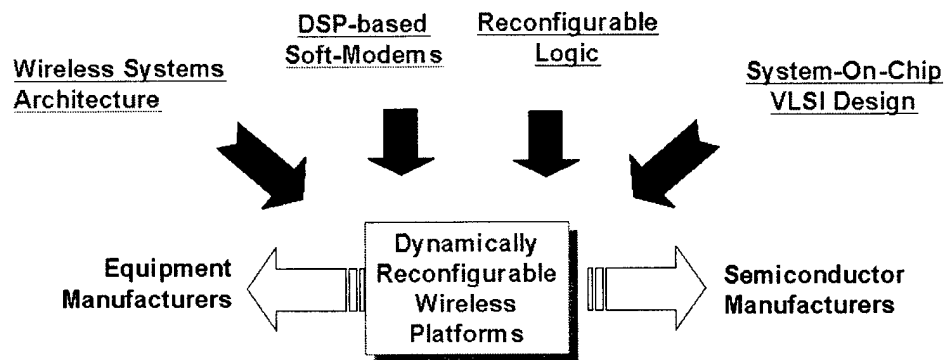
Degree of Resource Sharing, Granularity, and Operations per Data Sample:

A fundamental insight into the trade-off of architecture parameters is the strong correlation between the desirable degree of resource sharing, the chosen granularity of the macro- and micro-architecture, and the number of required operations per data sample:
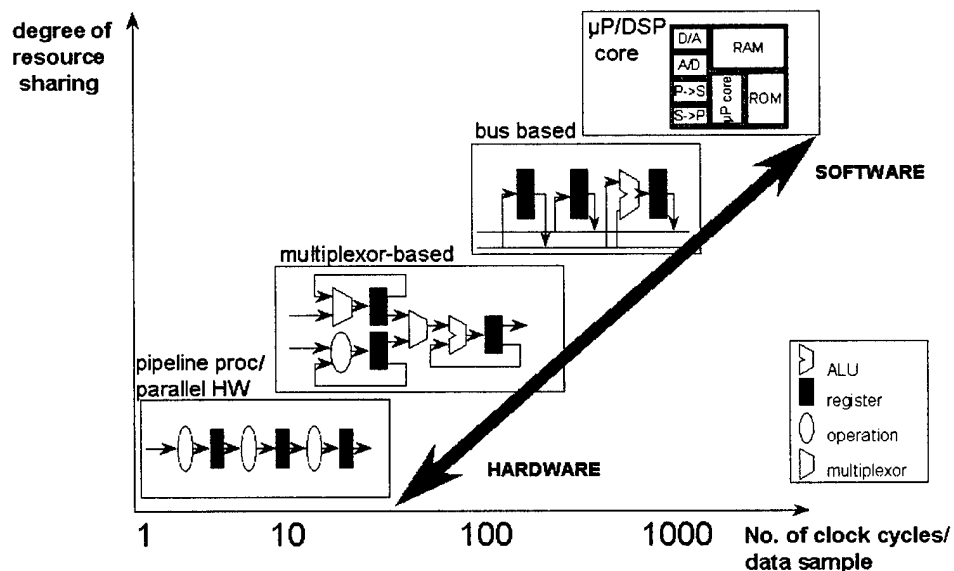
Granularity:

The following components lend themselves to conversion to ASIC cores:
Equalizers, Viterbi Decoders, Filters, CDMA Rake Receiver fingers

Technology Limits:

Present and soon to be released DSP devices are capable of handling GSM, IS136, and PDC however at significant power cost (poor battery life). For all the emerging high-speed data modes (IS136, EDGE) it appears that a significant part of the processing is too fast for DSP implementation. Present and future CDMA standards require substantial parts of the processing to be performed in hardware.

FPGA implementation appears to provide sufficient speed for all non-DSP processing at baseband for existing and foreseen services. Possibly the most compute intensive operation is the Viterbi Decoding for high-speed data services. Initial indications are that this can be achieved in present FPGAs

Physical Optimization:

Using a combination of ASIC cores, FPGA, and (DSP and/or RISC processor) appears to be the optimal solution. For CDMA the DSP is distributed across several embedded devices (QualComm, the leading US CDMA supplier uses an Intel 186 core). For TDMA and possibly the emerging WCDMA standards an augmented RISC processor with external reconfigurable logic for high-speed DSP algorithms may prove to be the optimal solution.

Chip Architecture:

It appears that all the Baseband processing can be implemented in one ASIC with FPGA structures included.

Mapping Algorithms:

For the GSM TDMA standard studied the most compute intensive algorithms, (MLSE Equalizer and Convolutional Decoder were implemented in FPGA.

Algorithmic Complexity, Moore's Law, and Power:

Functional complexity of wireless protocols is increasing faster than Moore's Law, while battery capacity increases only 2% per year! Therefore, the *energy per computation* has to be increasingly reduced over time by improving the system architecture at all levels, from low power circuits to well matched instruction sets, all the way up to faster algorithms.

Morphics Hardware Testbed:

Morphics hardware testbed is based on the AT&T / Lucent Technologies GSM test platform, Sceptre. The PWB contains a Xilinx part (implementing an accelerator), a crystal oscillator, and SRAMs. The Xilinx part is a memory mapped device relative to the Lucent 1618 DSP. Flat ribbon cables carries the DSP Data Bus, Address Bus and Control signals between the DSP and the accelerator. This testbed was used to demonstrate the MLSE Equalizer.

Morphics technology insertion into Lucent's Sceptre test platform:

Morphics flexible hardware test platform makes use of DSP processors (DSP1618, or ARM), low power controllers (CSP1088) and memory-mapped coprocessors (Xilinx FPGA) on daughter cards. All the resources on this platform can be accessed with reconfiguration if needed by a particular service implementation.

Four groups of services (PDC, IS-54 & IS-136, GSM, DCS1800 & PCS1900) that are now implemented as four separate boards, will be mapped onto one architecture running on the Morphics hardware test platform. The TDMA section is operational now:

Morphics Conclusion for this DARPA/Stanford Contract:

The Design completed and demonstrated on the Hardware Platform, (the most compute intensive parts of the GSM protocol) together with the analysis of IS-136 (North American TDMA), and IS-95 (CDMA), indicate that a multi-mode communication device is possible using reconfigurable architecture in combination with instruction-set processor(s) for baseband digital processing.

The Morphics development at time of conclusion of the contract successfully shows that a complete range of TDMA standards can be implemented. The next stage would be to add the CDMA capabilities.

Morphics Commercial Implementation beyond the DARPA/Stanford Contract:

This work can be considered as technology-transfer.

384 kb/s Convolutional Decoder:
Constraint length 9; Programmable Rates: 1/2, 1/3, 1/4; Programmable: Polynomials.
Status: Design and simulation complete, hardware tested.

IS-95 Rake Finger
Search capable; De-rotation; F & T tracking; PN de-spreading; Walsh channel decoding.
Status: hardware tested.

Reed Solomon Decoder
Codes 240, 192 / $GF(2^8)$; Input to 500 kb/s; Programmable: GF Polynomial Progr. n, k
Status, designed: syndrome calc., error polynomial gen., & solver; next: error correction.

Turbo Code Decoder
Block sizes to 1146 Iterations to 10.
Status: in design.

4.2.2 enVia Final Report

**Re-configurable Multi-mode, Multi-band Information Transfer Systems
RF Multiple-Service Front-end**

Summary

enVia's final report is contained in the Appendix C.

The **enVia** part of the project focused on the RF front end of a multiple service prototype system would support GSM and TDMA & AMPS operating at cellular, 800-900 MHz, and PCS, 1800-1900 MHz frequencies. An RF front end has been demonstrated that supports 4 modes and 2 bands, the cellular bands for AMPS, IS-136 as well as IS54A, IS54B, IS54C, and IS-95, as well as the PCS band, PCS1900 (upbanded GSM), IS136+ (upbanded IS-136), and IS-95+ (upbanded IS-95).

RF hardware considerations, granularity: LSI scale granularity discrete analog components are current best solution. VLSI scale single chip solutions are available but are just beginning to exhibit acceptable performance. Components that remain as discretes: SAW filters, Dielectric Resonator Filters, Crystals.

Digital /RF Interfaces: The project implemented an interface between the final Demod stage in the RF Front End and the high speed signal processing stage (reconfigurable logic stage). The interface signals from the RF Front end include: Analog, Baseband, and I&Q.
Control signals coming into the RF Front End are also analog. In order to support the full range of services, *three* different *interface bandwidths* are provided.

Below, we highlight parts of the enVia presentation-style final report.

enVia Project Summary:

The project components have been created and developed to a demonstrable level as per the project plan:

1.	The Digital and RF sections have been demonstrated.

2.	Each section works independently according to the project plan, with the exception that the digital section does not support all the modes planned.

3.	The interoperability of the systems can be implied from testing done separately (could be demonstrated with additional funding.)

Overview:

The Project Objectives Delivered:

Prototypes demonstrated with RF Front End supports *4 modes and 2 bands* :
Cellular Band, AMPS, IS-136 as well as IS54A, IS54B, IS54C, IS-95, PCS Band, PCS1900 (upbanded GSM), IS136+ (upbanded IS-136), IS-95+ (upbanded IS-95).

Function Mapping:

All of these services can be decomposed in the following basic functions:

RF Front End,
High Speed Signal Processing,
Low Speed Signal Processing,
Control.

The best distribution of these functions considering available technology at this time is:

RF Front End — analog discrete components,
High Speed Signal Processing - FPGA,
Low Speed Signal Processing - DSP,
Control — Microcontroller.

Reconfigurable RF Architecture:

Analog technology is currently the best way today of implementing a
multi-mode multi-band RF frontend.

In order to achieve the desired performance, this study concluded that
currently a design based on discrete components meets the goals, esp. cost.

Chip set solutions are continuing to improve and
will soon have adequate performance.

RF Granularity:

LSI scale granularity discrete analog components are current best solution

VLSI scale single chip solutions are available but are
just beginning to exhibit acceptable performance

Components that remain as discretes
SAW filters
Dielectric Resonator Filters
Crystals

Sample RF Front End Design supporting *4 modes, 2 bands, and 3 bandwidths* :

RF Performance:

- Receiver *(1.23 MHz IF Bandwidth)*

    - Frequency Range:
        – LoBand: *869 to 894 MHz*
        – HiBand: *1930 to 1990 MHz*

    - Sensitivity:
        – LoBand: *-88 dBm (Receiver compressed by off-the-air signals)*
        – HiBand: *-103 dBm*

    - Input 3rd Order Intercept
        – LoBand: *-5 dBm*
        – HiBand: *-12 dBm*

    - AGC Range: *>75 dB*

    - Demod output: *Analog I and Q signals, balanced, 1Vpp*

    - DC Current (Battery Voltage = 3.6 V):
        – LoBand: *66 mA*
        – HiBand: *76 mA*

- Transmitter

    - Frequency Range
        – LoBand: *824 to 849 MHz*
        – HiBand: *1850 to 1910 MHz*
        – Output: *-3 dBm ___2dB*
        – AGC Range: *> 64 dB*

    - Demod input: *Analog I and Q signals, balanced, 1Vpp*

    - DC Current (Battery Voltage = 3.6 V) : *78 mA*

RF Response, LoBand:

RF Response, HiBand:


```
                                                    25Feb1998  19 06 23
CH1   S21    logMAG      10dB/     REF 0 dB        1_-50.826dB
      b.aPCS R X                                   1 850.00  040 MHz
PRm                                                2_ 25 598 dB
Cor                                                   1.91 GHz
                                                   3_ 8 25 46 dB
Avg                                                   1.93 GHz
16                                                 4_ 8 62 11 dB
                                                      1 99 GHz
                                                   5_ 8 68 63 dB
                                                      1.954 GHz



START 1 800.000 000 MHz              STOP 2 000.000 000 MHz
```

RF Design Conclusions:

- The initial design concept was to use highly integrated RF functions. However, as the project matured, a larger mix of discreet components were used than originally expected in order to meet project goals for performance, power efficiency, size, and cost.

- To meet the requirements for IS-136, GSM, and IS-95 services, three IF SAW filters were used with center frequencies chosen from commercially available units with frequencies close to 110 MHz. In the latest version of the design, the IF frequency has been increased to approximately 211 MHz for a reduction in package size and cost.

- The final version of the design is anticipated to meet objectives for performance, size, battery life, and cost.

Digital /RF Interfaces:

The project implemented an interface between the final Demod stage in the RF Front End and the high speed signal processing stage (reconfigurable logic stage).

The Information interface from the RF Front end contains: Analog, Baseband, I&Q.
Control signals coming into the RF Front End are also analog.
To support the full range of services *3 Interface bandwidths* are provided.

RF Physical Optimization:

Current RF design: can be shrunk, chip set size estimate is 2.8 x 1.5".

RF Chip Architecture:

A currently available chip set solution (not prototyped in this project) has:
4 chips and requires *several 100* off chip components.

RF Test Bed:

Test Bed *prototypes* demonstrated RF Front End, supports *4 modes and 2 bands* :
   Bands: Cellular, AMPS, IS-136 as well as IS54A, IS54B, IS54C, IS-95, PCS,
   PCS1900 (upbanded GSM), IS136+ (upbanded IS-136), IS-95+ (upbanded IS-95).
The RF Brass Board prototype is composed of five test boards,
      it demonstrates all currently deployed US cellular and PCS services.

5.1 Personnel

Stanford University:

Principal Investigator: Michael J. Flynn
Address: E.E. Dept., Gates Computer Science Building, Room 334, Stanford, CA 94305,
email: flynn@ee.Stanford.edu, Phone: (650) 723-1450, Fax: (650) 725-6949

Investigator: Martin Morf
Address: E.E. Dept., Gates Computer Science Building, Room 335, Stanford, CA 94305,
email: morf@arithmetic.Stanford.edu, Phone: (650) 723-0140, Fax: (650) 725-6949

Investigator for security:
        Made major contributions to this project, but was not supported by this contract.
        John T. Gill III.

Supported Graduate Students:
        Oskar Mencer,   (PamBLOX, Reconf. Alg. Arch., & Arithmetic.)
        Hon-Cheong Leung,   (PamBLOX designs, DES, IDEA, with O. Mencer.)
        Kevin W. Rudd,   (VLIW Architectures, RF Link operations.)
        Patrick S.-Y. Hung   (System Admin., VLSI Floor-Planning.)
        Hyuk-Jun Lee,   (Reconfigurable Arithmetic)
        Alice C. Yu,   (Efficient Video Encoding)
        Andrew C. Zimmerman   (Parallel Architectures and Simulation Tools.)

Other Graduate Students:
        Made contributions to this project, but were unsupported by this contract.
        Jorge Campello De Sousa,   (Communication and Security algorithms, with J. Gill)
        Craig Dowie,   (Stanford Security Demo)
        Luc R. Semeria,   (Logic Synthesis, with Oskar Mencer.)
        Hossam A. H. Fahmy   (Low-Power Quantum -Dot Logic, with M. Morf.)

Academic Visitor:
        Made contributions to this project, but were unsupported by this contract.
        J.-M. Delosme, Evry Univ., FR.(Algorithms and Logic, with M. Morf.)
        A. Huang, Terabit   (Network Architectures, Photonics, with M. Morf.)
        M. T. Hadidi, Mobil   (Algorithms for distributed Systems, with M. Morf.)

enVia team:

Subcontractor: Mark Cummings, enVia Inc.,
Address: 348 Camino al Lago Atherton, CA. 94027, e-mail: cummings@envia.com,
Phone/Fax: (650) 854-4406

Contact enVia: Ken Jacobsen, COO enVia Inc.
Phone: +1.408.777-4804, Fax: +1.408.873.0336, e-mail: kjacobsen@envia.com

RF design, enVia: Rich Walsworth
Engineer, enVia   John Ralston
DSP design, Cons.   Steve Sweitzer

Morphics team:

Contact Morphics   Colin McNab,   e-mail: macnab@morphics.com

Engineer, Morphics   David Holmes
Engineer, Morphics   Stephen Wasson

5.2 Publications

[CGMF Nov97]   J. Campello, J. T. Gill, M. Morf, M. J. Flynn,
Smart Photonic Networks and Computer Security for Image Data,
SPIE International Symposium on Voice, Video, and Data Communications,
Dallas Texas, Nov. 1997. (invited)

[FKM Sep98]   H. A. H. Fahmy, R. Kiehl, M. Morf,
Quaternary Phase Quantum -Dot Logic,
IEEE Transactions on Computers, submitted Sep. 1998

[FR CRC97]   M. J. Flynn and K. W. Rudd,
Parallel Architectures
The Computer Science and Engineering Handbook, CRC Press, 1997.

[FR Mar96]   M. J. Flynn and K. W. Rudd,
Parallel Architectures,
ACM Computer Surveys, March 1996.

[MF Jul99]   O. Mencer, M. Morf, M. J. Flynn,
Precision of Semi-Exact Redundant Continued Fraction Arithmetic for VLSI,
SPIE '99 (Arithmetic session), Denver, July 1999.

[MF Oct98]   M. Morf, M. Flynn,
Architectural Issues and Nano-Technology,
special session on Photonics and Nano-Technology,   Chair L.S. Lome, BMDO,
Optical Society of America Conference, Baltimore, Oct. 1998. (invited)

[MHA... Dec97]   W.H. Mangione-Smith, B. Hutchings, D. Andrews, A. DeHon, C. Ebeling,
R. Hartenstein, O. Mencer, J. Morris, K. Palem, C. V. Prasanna, H. Spaanenburg,
Seeking Solutions in Configurable Computing ,
IEEE Computer Magazine, Dec. 1997.

[MP Jan99]   O. Mencer, M. Platzner,
Dynamic Circuit Generation for Boolean Satisfiability in an Object-Oriented Design Environment,
Hawaii International Conference on System Sciences (ConfigWare Track), Jan. 1999

[MMDF Nov98]   M. Morf, O. Mencer, J.-M. Delosme, M. J. Flynn,
Reconfigurable Computing and CORDIC Like Architectures,
ICCU'98, Korea, November , 1998. (invited)

[MMF sub99] O. Mencer, M. Morf, M. J. Flynn,
Programmability, Performance, and Power of FPGA based Computing Machines,
(submitted for publication to Transec99)

[MMF Jul99] O. Mencer, M. Morf, M. J. Flynn,
Precision of Semi-Exact Redundant Continued Fraction Arithmetic for VLSI,
SPIE '99 (Arithmetic session), Denver, July 1999.

[MMF Apr98] O. Mencer, M. Morf, M. J. Flynn,
PAM-Blox: High Performance FPGA Design for Adaptive Computing,
IEEE Symposium on FPGAs for Custom Computing Machines, Napa Valley, April 1998.

[MMF May98] O. Mencer, M. Morf, M. J. Flynn,
Hardware Software Tri-design of Encryption for Mobile Communication Units,
IEEE International Conference on Acoustics, Speech and Signal Processing, May 1998.

[MMF Jun98]   O. Mencer, M. Morf, M. J. Flynn,
Pipelined CORDICs for Reconfigurable Computing,
The Sixth FPGA/PLD Design Conference Exhibit,   Pacifico Yokohama, Yokohama, Japan, June 24-26, 1998.

[MSDM Nov98] O. Mencer, L. R. Semeria, J.-M. Delosme, M. Morf,
Application of Reconfigurable CORDIC Architectures,
Asilomar Conference on Signals, Systems, and Computers, California, Nov. 1998.

[MSF Fal98]   O. Mencer, M. Shand, M. J. Flynn,
FireLink: A High-Performance Adaptive Firewire Interface,
IEEE Computer, special issue on High Performance Network Interfaces, Fall 1998.   (Digital Systems Research Center TechNote 1998-012).

[RF Mar97]   K. W. Rudd and M. J. Flynn,
Instruction-level Parallel Processors---Dynamic And Static Scheduling Tradeoffs.
Proceedings of IEEE International Symposium on Parallel Algorithms Architecture Synthesis, March 1997.

[WM Apr99]   F. Whartman, M. Morf,
The Past and Future of Architecture,
Microprocessor Workshop, Asilomar, April 21-23, 1999.

[Y Dec98]   A. Yu,
Low Complexity Video-Encoding and Quality Metrics
Stanford University, E.E. Dept. Ph.D. Thesis, to be submitted Dec., 1999.

[YF Nov98]   A. Yu and M. Flynn,
Implementation and Optimization Issues for the H.263 Compression Standard,
Proceedings 32nd Asilomar Conference on Signals, Systems, and Computers, Nov. 1998.

[YLF Feb97]   Alice Yu, R. Lee, and Michael Flynn,
An Evaluation of Video Fidelity Metrics,
COMPCON Digest of Papers, San Jose, California, pp. 49-60, Feb., 1997.

[YLF Sep97]   Alice Yu, Ruby Lee, and M. J. Flynn,
Early Detection of All-Zero Coefficients in H.263,
Proceedings of the Picture Coding Symposium, Berlin, Germany, pp. 159-164, Sep., 1997.

[YLF Nov97]   Alice Yu, Ruby Lee, and Michael Flynn,
Performance Enhancement of H.263 Encoder Based on Zero Coefficient Prediction,
Proceedings of the Fifth ACM International Multimedia Conference, Seattle, Washington, pp.21-29, Nov. 1997.

Theses, CSL Technical Reports, Presentations (ppt files):

[D Feb98]   Craig Dowie, presentation as part of the Stanford Security Demo to Darpa in February 1998.
http://umunhum.stanford.edu/res html/darpa/com95-97.ppt

[L 98]   Hon-Cheong Leung, Project Report on PamBLOX designs for DES and IDEA Encryption, with O. Mencer, 1998.

[LF Feb99] H.-J. Lee and M. J. Flynn
Coarse Grain Carry Architecture for FPGA.
Stanford University, Feb. 1999. [CSL-TR-99-780]

[LF Oct98]   H.-J. Lee and M. J. Flynn
Designing a Partitionable Multiplier.
Stanford University, Oct. 1998, [CSL-TR-98-772]

[M Dec99]   O. Mencer, Stanford University, E.E. Dept. Ph.D. Thesis, Dec., 1999.

[MF TR97]   O. Mencer, M. J. Flynn,
A Selection of Recent Advances in Computer Architecture,
Stanford, Computer Systems Laboratory, Technical Report, 1997.[CSL-TR-97-745]

[R Dec99]   K. W. Rudd, Stanford University, E.E. Dept. Ph.D. Thesis, Dec., 1999.

[R May98]   K. W. Rudd,
VLIW Processors---Efficiently Exploiting Instruction-level Parallelism.
Stanford University, May 1998, Ph.D. Defense.

[R Apr98]   K. W. Rudd,
Replay Buffers---Effective Dynamic Event Management for VLIW Processors,
Twenty-Fourth Annual Asilomar Microcomputer Workshop, April 1998

[R Apr98]   K. W. Rudd,
Merced: Gateway to Yosemite.
Twenty-Fourth Annual Asilomar Microcomputer Workshop, April 1998, Panel session.

[R Aug97]   K. W. Rudd,
Efficient Exception Handling for High-performance Processor Architectures.
Stanford University, August 1997, [CSL-TR-97-732].

5.3. Project URL's

Stanford University, Algorithms and Architecture Group, Prof. M. Flynn, PI:

http://umunhum.stanford.edu/
http://umunhum.stanford.edu/res_html/darpa/radio.html
http://umunhum.stanford.edu/PAM-Blox/
http://umunhum.stanford.edu/~oskar/pubsMedium.html

Morphics Inc.:

www.morphics.com

6. Demos

Summary of Project Goals:

The goals of this project were:

1.   Services:
     Sample prototype modes: AMPS, PCS-1900; 800-900, 1800-1950 MHz.

2.   Service reconfiguration time (mode and/or band reconfiguration) prototype 800 ms.
     Actual simultaneous modes are not expected to be supported, except for plans
     to use ASICS for instance for paging that could operate simultaneously.
     Service reconfiguration time goals are set to allow duplexing or switching
     between the services without service interruption (taking advantage of long
     timeouts, of the order of a second.)

3.   Security Prototype
     Single Path FPGA implementation of encryption.
     Two paths in software.
     Stanford to demonstrate the goals on a Stanford testbed.

Demos:

In Summary:

1. Stanford demo, as proposed, see also Appendix A..
2. Morphics demo, TDM running, unable to deliver GSM, see also Appendix B.
3. enVia demo, designs complete, 5 working boards, final version late, see Appendix C.

6.1. Stanford demo:

Stanford presented in February 1998 a Security Prototype demo as proposed, consisting of:

A Single Path FPGA implementation of encryption (IDEA, DES) in software and
downloadable to FPGAs on DEC Pamette board.

A Two Path security demo in software, running on three laptops, February 1998.
Graduate Studenty presentation by Craig Dowie (PowerPoint file).
Types of links supported:
wireless (Ricochet, Freewave), modem, ethernet, IR, serial.

Summary of the Stanford Demo presentation to Darpa (see Appendix C.):

Project Goals of Demo:

Investigation into improving the *security and reliability* of data transmissions
between hosts by the use of multiple, physically disjoint communications links:

Naturally Disjoint Wireless Links (e.g. RF modems, Infra-Red, etc),
Distinct Network Links on WAN (e.g. Internet)

Demo Concepts, Simple Communication & Security Example:

Demonstrates some concepts behind project work: .Simple Secret Sharing

Demonstrates Secure Transmission of Data between Hosts
using Two Physically Disjoint Paths.

Illustrates how a Single Path can be Compromised while Maintaining Data Security.
Uses a trivial (XOR) encryption scheme (one instruction per data word).

Demo System Components:

Hosts: 3 Pentium MMX Laptops with additional serial comms interfaces (PC card)

Operating System Environment:: Windows NT 4.0   (not a critical choice)

Connections:
  Cable with Null modems,
 Ricochet SX RF wireless modems;
Additional Possible Paths (not used in demo):
  FreeWave DGR-115/H RF wireless modems,
  Ethernet (WAN, Internet),
  Modem (e.g. Phone)
  Infra-Red.

Demo Physical Configuration:

3 physical connections

Arapahoe to Cree :	Wireless Ricochet
Sioux to Cree :	Wireless Ricochet
Arapahoe to Sioux :	Null Modem (RS-232)

```
┌─────────────────────────────────────┐
│   Physical Configuration            │
└─────────────────────────────────────┘
```

```
              ┌──────────┐
              │ Ricochet │
              │   Cree   │
              │ Ricochet │
              └──────────┘
             ↗            ↖
           ╱                ╲
         ╱                    ╲
   ┌──────────┐          ┌──────────┐
   │ Ricohet  │          │ Ricochet │
   │ Arapahoe │◄────────►│  Sioux   │
   │  Null    │          │  Null    │
   └──────────┘          └──────────┘
```

Network Configuration (TCP/IP):

The connections are designed to model two disjoint data paths:
  between hosts Arapahoe and Sioux (the left and right laptops)
Cree (in the middle) models an eavesdropper -
  'snooping' on one of the paths (or TTP)

In practice, Cree stores from Arapahoe and forwards it to Sioux.

Demo Information Flow:

Arapahoe: reads data to be sent from a file; generates a key and encrypts the data.
Transmits the data along null modem connection (directly to Sioux).
Transmits the key along the wireless (Ricochet) path to Sioux (via Cree).
Cree 'intercepts' and displays the key .

Demo Conclusion and Continuing Work:

Multiple path socket class based on TCP/IP sockets.
Multiple path and node management software (COMPASS)
Further investigation of wireless and communications hardware
    (e.g. wireless ethernet technology, digital radio, Ultra-Wide Band technologies)

The required system integration work far outweighs the work required to implement our security concepts. The required system integration work is much more generally useful, i.e. for improving performance and reliability by enabling parallelism. We anticipate that *future parallel operating systems* will *support our security and reliability concepts* much more readily!

6.2 enVia Final Demos

enVia was responsible for research and development of a reconfigurable hardware test bed aimed at deployed mobile wireless services. Late into the contract enVia was reorganized, but the reorganized enVia/Morphic was transparent to the research contract, as all work was performed under enVia's subcontract; however, the demos in this final report were organized recognizing this change of responsibility.

6.2.1 Morphics Multi-Standard Digital Hardware Demo Summary

The Morphics part of the project initial intent was to show that multiple digital standards can be implemented in a single baseband architecture without multiple, parallel processing paths. This development would then be incorporated with a separate multi-band RF development to create a single multi-standard, multi-band communications device

The scope of the project as originally conceived proved too big to attempt in its entirety. An end-to-end mobile handset is by its nature a $50M project. An attempt to use commercially available platforms was partially implemented in that the development hardware was obtained. However, no software was available, so a complete multi-standard baseband implementation was not attempted.

The *target development* then became a multi-standard hardware logic section, the only new, previously undeveloped part of a multi-standard communications baseband, to enable a *multi-standard digital modem.*

The initial candidate services were AMPS, GSM, IS136, and IS95, i.e. analog cellular, European TDMA, North American TDMA, and CDMA.

The services considered during the scope of the project were all of the TDMA based services: GSM: 800, 1800, 1900; IS-54 & IS-136; PDC.

Two part Morphics Demo:

        a) GSM FPGA baseband demo.
        b) Software structural design for second service.

The GSM standards were implemented and the IS-54/136 were in process at time of conclusion of the project. The CDMA services would have been the next logical development under an extension to the existing program.

A test bed was assembled based on the Lucent Technologies GSM test platform (Sceptre). A PWB containing a Xilinx part (implementing an accelerator), a crystal oscillator, and SRAMs was designed and fabricated. This made the Xilinx part a memory mapped device relative to the Lucent 1618 DSP. Flat ribbon cables carried the DSP Data Bus, Address Bus and Control signals between the DSP and the accelerator. This was used to demonstrate the MLSE Equalizer.

Morphics Demo conclusion: The Design completed and demonstrated on the Hardware Platform, (the most compute intensive parts of the GSM protocol) together with the analysis of IS-136 (North American TDMA), and IS-95 (CDMA), indicate that a multi-mode communication device is possible using reconfigurable architecture in combination with instruction-set processor(s) for baseband digital processing. The Morphics development at time of conclusion of the contract successfully shows that a complete range of TDMA standards can be implemented. The next stage will be to add the CDMA capabilities (beyond this contract).

Digital Technology Limits: Present and soon to be released DSP devices are capable of handling GSM, IS136, and PDC however at significant power cost (poor battery life). For all the emerging high speed data modes (IS136, EDGE) it appears that a significant part of the processing is too fast for DSP implementation. Present and future CDMA standards require substantial parts of the processing to be performed in hardware.
Some combination of ASIC cores, FPGAs, DSP and RISC cores appears to generate well matched solutions. FPGAs are well suited and flexible for wireless communication; however, a tri-codesign (software-, firmware-, and hardware-codesign) can significantly improve area and power efficiency, while optimizing flexibility. Morphics for instance has found that they can reduce the silicon area by a factor of 45 with structured ASIC designs over FPGAs. The focus of their work now is the design of "silicon IP" -- cores.
For CDMA the DSP is distributed across several embedded devices. QualComm, CDMA supplier, uses an Intel 186 core. For TDMA and possibly for WCDMA, an augmented RISC core with reconfigurable logic for high speed DSP algorithms appears to be a matched solution, i.e. combining controller and DSP.

6.2.2 enVia Multiple-Service RF Front-end Demo Summary

After its reorganization *enVia's* work focused on the RF front end of a multiple service prototype system, that would support GSM and TDMA & AMPS operating at cellular, 800-900 MHz, and PCS, 1800-1900 MHz frequencies. An RF front end has been demonstrated that supports 4 modes and 2 bands, the cellular bands for AMPS, IS-136 as well as IS54A, IS54B, IS54C, and IS-95, as well as the PCS band, PCS1900 (upbanded GSM), IS136+ (upbanded IS-136), and IS-95+ (upbanded IS-95).

RF hardware considerations, granularity: LSI scale granularity discrete analog components are current best solution. VLSI scale single chip solutions are available but are just beginning to exhibit acceptable performance. Components that remain as discretes: SAW filters, Dielectric Resonator Filters, Crystals.

Digital /RF Interfaces: The project implemented an interface between the final Demod stage in the RF Front End and the high speed signal processing stage (reconfigurable logic stage). The interface signals from the RF Front end include: Analog ,Baseband, I&Q.

Control signals coming into the RF Front End are also analog. In order to support the full range of services, *three* different *interface bandwidths* are provided.

RF Demo: RF System consists of the following completed five Test Boards: Receiver Test Board, Transmitter Test Board, 1st LO Test Board, IF Test Board, Combiner Test Board. Switchboard and Auxiliary LO

RF Design Conclusions: The initial design concept was to use highly integrated RF functions. However, as the project matured, a larger mix of discreet components were used than originally expected in order to meet project goals for performance, power efficiency, size, and cost.

To meet the requirements for IS-136, GSM, and IS-95 services, three IF SAW filters were used with center frequencies chosen from commercially available units with frequencies close to 110 MHz. In the latest version of the design, the IF frequency has been increased to approximately 211 MHz for a reduction in package size and cost. The final version of the design is anticipated to meet objectives for performance, size, battery life, and cost.

RF Technology Limits: Current RF design is capable of supporting all target services. Limitations in the current design are in size, and power consumption: Current size limit is 2.8" by 1.5" for chip set solution not prototyped in this project, can be shrunk further.

7. Deliverables

Deliverables on this contract include:

1.   Reports and Presentations:

Darpa Quarterly reports,

Darpa PI meeting presentations,

Darpa site visit presentations.

2.   Stanford Darpa project website:

http://umunhum.stanford.edu/res_html/darpa/radio.html
http://umunhum.stanford.edu/PAM-Blox/

http://umunhum.stanford.edu/~oskar/pubsMedium.html

3   Morphics Inc.:

www.morphics.com

4.   Demos, on-site:

Stanford security demo, February 1998.

enVia demos:

Digital IF / Baseband by Morphics

RF Frontend by enVia

5.   Final Report:

Stanford (Word and PowerPoint file, see also website above)

enVia Final Report:

enVia   (Word, and PowerPoint file)

Morphics   (PowerPoint file)

8. Technology Transfer

Stanford interacted with J. Massey from the ETHZ and Cylink on security algorithms (e.g. IDEA). His latest versions (FASTER) seems to be better adapted to FPGA technologies.

Morphics has several major industrial and financial partners
(e.g. AT&T / Lucent SCEPTRE board)
        www.morphics.com website lists:
                July 20, 1999, Morphics closes second round financing of $13.5million.

REPORT DOCUMENTATION PAGE

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	May 17, 2000	FINAL REPORT 09-06-96 – 09-05-99

4. TITLE AND SUBTITLE	5. FUNDING NUMBERS
Reconfigurable Multimode, Multiband Information	C-DABT63-96-C-0106-P00005 PR- SPO # 17231 Fund # 182A015 Acct. # 2DSA707

6. AUTHOR(S)

Michael J. Flynn

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
COMPUTER SYSTEMS LABORATORY STANFORD UNIVERSITY STANFORD, CA 94305	

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSORING / MONITORING AGENCY REPORT NUMBER
ONRRO UNIVERSITY OF WASHINGTON 1107 NE 45TH STREET, SUITE 350 SEATTLE, WA 98105-4631	

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION / AVAILABILITY STATEMENT	12b. DISTRIBUTION CODE
Approved for public release: distribution unlimited.	

13. ABSTRACT (Maximum 200 words)

14. SUBJECT TERMS

15. NUMBER OF PAGES
46 pages

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
unclassified	unclassified	unclassified	unlimited

NSN 7540-01-280-5500